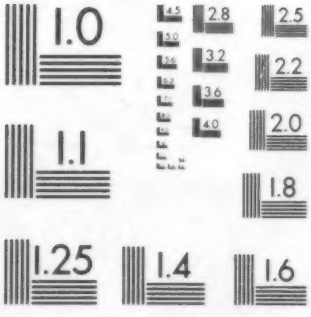

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS
STANDARD REFERENCE MATERIAL 1010a
(ANSI and ISO TEST CHART No. 2)

DOCUMENT RESUME

ED 230 186 IR 010 707

AUTHOR Hayes-Roth, Barbara
TITLE Implications of Human Pattern Processing for the Design of Artificial Knowledge Systems.
INSTITUTION Rand Corp., Santa Monica, Calif.
REPORT NO Rand-P-5854
PUB DATE Apr 77
NOTE 31p.
AVAILABLE FROM Rand Corporation, Main St., Santa Monica, CA 90406 ($3.00).
PUB TYPE Information Analyses (070)

EDRS PRICE MF01 Plus Postage. PC Not Available from EDRS.
DESCRIPTORS *Artificial Intelligence; *Cognitive Processes; Cues; *Cybernetics; *Design Requirements; Epistemology; Literature Reviews; *Pattern Recognition
IDENTIFIERS *Pattern Matching; Pattern Meaning; Prototype Acquisition; Stimulus Complexity

ABSTRACT

This paper presents evidence that four design principles commonly embodied in artificial knowledge systems are inconsistent with human cognitive capabilities. Because these principles are widely accepted as characteristics of human knowledge processing, common theoretical properties related to cognitive psychology and artificial intelligence which are apparent in pattern processing research, are considered, including (1) canonical meaning representations of stored patterns; (2) prototypes as representations of classes of stored patterns; (3) pattern-matching processes that must fully evaluate all stored patterns that partially match a target pattern; and (4) pattern-matching processes that are influenced by pattern complexity and limited processing capacity. This characterization of human pattern processing is critically evaluated in light of recent psychological studies, and the implications of these results for artificial knowledge systems design are discussed. Forty-two references are listed. (LMM)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

U.S. DEPARTMENT OF EDUCATION
NATIONAL INSTITUTE OF EDUCATION
EDUCATIONAL RESOURCES INFORMATION
CENTER (ERIC)

- This document has been reproduced as received from the person or organization originating it.
- Minor changes have been made to improve reproduction quality.

- Points of view or opinions stated in this document do not necessarily represent official NIE position or policy.

ED230186

IR

IMPLICATIONS OF HUMAN PATTERN PROCESSING FOR THE DESIGN OF ARTIFICIAL KNOWLEDGE SYSTEMS

Barbara Hayes-Roth

Rand Corp. Santa Monica

April 1977

"PERMISSION TO REPRODUCE THIS MATERIAL IN MICROFICHE ONLY HAS BEEN GRANTED BY

Rand Corp.

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

P-5854

3.00

IR010707

The Rand Paper Series

Papers are issued by The Rand Corporation as a service to its professional staff. Their purpose is to facilitate the exchange of ideas among those who share the author's research interests; Papers are not reports prepared in fulfillment of Rand's contracts or grants. Views expressed in a Paper are the author's own, and are not necessarily shared by Rand or its research sponsors.

The Rand Corporation
Santa Monica, California 90406

Abstract

Artificial knowledge systems typically embody several design principles that are also widely accepted as characteristic of human knowledge processing. These properties include: canonical meaning representations for stored patterns; prototypes as representations of classes of stored patterns; pattern matching processes that must evaluate fully all stored patterns that partially match a target pattern; and pattern matching processes that are influenced by pattern complexity and limited processing capacity. This characterization of human pattern processing is critically evaluated in light of recent psychological studies. The implications of these results for the design of artificial knowledge systems are discussed.

Introduction

Parallel theoretical developments in cognitive psychology and artificial intelligence are readily apparent in research in pattern processing, that is, in theories of pattern representation and matching and pattern-directed inferencing. Several common theoretical properties are considered in the present paper. These include: canonical meaning representations of stored patterns; prototypes as representations of classes of stored patterns; pattern-matching processes that must evaluate fully all stored patterns that partially match a target pattern; and pattern-matching processes that are influenced by pattern complexity and limited processing capacity. The first two properties are theoretically motivated by presumed efficiencies of representation, matching, and inferencing. The second two properties are assumed to be inevitable properties of all pattern processing systems, even though they impair performance by increasing pattern matching times and errors. Despite the wide acceptance of the theoretical assumptions introduced above, recent studies suggest that they are inconsistent with human pattern processing capabilities. The assumptions are critically evaluated in light of these results below. The implications of these results for the design of artificial knowledge systems are also discussed.

Canonical Meaning Representations of Stored Patterns

In canonical meaning representations [25, 38, 39, 42], the semantic constituents of patterns are abstracted and stored in standardized structures. Thus, the meanings, but not the surface features, of patterns are stored in memory. It has been argued that canonical meaning representations provide several efficiencies. They minimize storage requirements because semantic redundancies among patterns can be factored out to avoid redundancies among representations. For example, the proposition "Mary handed John the book and he accepted it" includes two references to the transfer of the book, "handed" and "accepted." Both of these might be replaced by a single semantic verb primitive such as "physical transfer." Canonical meaning representations also facilitate pattern matching by eliminating the need to search among semantically equivalent representations of propositions. As Schank [37] has argued:

> How should a proposition appear in memory? It simply cannot be stored in a form that bears a strong relationship to the original form of the input in natural language. Since it is possible to say the same thing in any number of different ways, it is unreasonable to suppose that people (or machines) are constantly checking to see if a proposition that they have stored one way in memory is the same as another they have stored somewhere else in a different form ...
> ... there must be a canonical form for meaning representations ...

Canonical meaning representations facilitate inferential reasoning in two ways. First, the representation of a

proposition makes explicit those meaning elements that support inferencing. For example, the proposition "John ran," implying movement of an agent x from location y to location z, would include a subrepresentation of the abstract inference, TRANSITION (x, y, z). Because the inference is represented explicitly in the meaning representation, inferential reasoning reduces from a deduction problem to a simple retrieval problem. Second, the use of canonical representations provides generality to inference rules that enable their application to a wider variety of situations. For example, because the propositions "Mary walked" and "Alice went" also imply movement of an agent x from location y to location z, their canonical representations would necessarily include subrepresentations of the same abstract inference, TRANSITION (x, y, z). As a result, the same general inferences would be entailed by all three propositions.

The assumption of canonical meaning representations is also consistent with prominent psychological theories, as illustrated in the following excerpt from Miller [30]:

> It is usually not the precise words we happen to use in characterizing an experience that are stored and later recalled. Rather, it is some prelinguistic, conceptual representation that we seem to remember--something nonverbal, but with an affinity for verbal expression, couched in the dimensions of verbal thought.

Several kinds of empirical evidence have been adduced to support the assumption that people have canonical meaning

representations. For example, people can recall studied information, given a recall cue that is a paraphrase of the verbatim cue present during study [4]. This suggests that the meanings of studied cues are stored in canonical representations. Any cue that preserves the meaning of a studied cue, regardless of its form of expression, will produce a successful match to the representation of the cue. People also frequently recall paraphrases of the material they have studied, rather than the literal wordings of the material [9]. This suggests that only meaning is retained from studied material. A meaning representation would predict recall of the gist, but not necessarily the exact wording of studied material. Finally, people frequently report that they recognize (have seen previously) paraphrases of material that they have seen previously [5, 6, 15, 16, 35]. This also suggests that the meanings, but not the surface features, of studied materials are stored in memory. Thus, any target that expresses the meaning of studied material will produce a successful match to its memory representation.

While these data are consistent with the assumption that people have canonical meaning representations, they do not exclude the possibility that people retain explicit information about the literal form of the material they acquire. In fact, it is possible to account for all of the data with a theory that assumes surface information, such as the specific words in a sentence, as the basic unit of

representation [[21]]. For example, consider the finding that people can recall material they have studied given a cue that paraphrases the originally presented cue. The theory assumes that people store literal representations of studied cues and associated material in an integrated semantic-propositional network. The theory also assumes that storage causes a spread of network "activation" to semantically related (including synonymous) information in the network and that activation creates a memory trace. In other words, this theory assumes that, during study, memory traces representing paraphrases of studied cues would be created along with memory traces representing the verbatim cues. These paraphrase cue traces would presumably enable subsequent recall of studied material, given a paraphrase cue. The other data described above can be accounted for in a similar manner.

Further, a reevaluation of the psychological data [21] supports the assumption of literal representations over the assumption of canonical representations. For example, although people do sometimes recall studied material given a paraphrase cue, they recall it more often when given a verbatim cue [4]. If people had canonical representations of the cues, then either one should activate the appropriate representation equally well. Therefore, recall should be equally good given either cue. The literal representation theory, on the other hand, assumes that direct activation of the verbatim cue representation during study gives it a stronger memory trace

than spreading activation gives the paraphrase cue representation. Therefore, it predicts that, as observed, the verbatim cue should be more effective. Similarly, although people do sometimes recall paraphrases of the material they have studied, they generally recall the literal form of the material at least as often as, and frequently more often than, any paraphrase [7]. If people had canonical representations of studied material, they should produce paraphrased expressions equally often during recall. The literal representation theory, on the other hand, again assumes that the literal form of the material has the stronger memory trace and so should be recalled better. Finally, although people frequently "recognize" material that paraphrases studied material, they recognize explicitly studied material more often and with greater confidence [35]. If people had canonical representations of studied material, any expression of the material should activate the appropriate representation. Therefore, studied and paraphrase forms should be recognized equally well. The literal representation theory again assumes that the literal form of the studied material has the stronger representation and so should be recognized better.

A recent study [21] provided additional support for the literal representation theory. In the first experiment, people read thirty-two simple sentences, such as "The spy tore up the documents." They were told to try to remember the meanings of the sentences. Later, they were given a surprise recognition test. The test included the originally read sentences and some

semantically equivalent sentences, in which the original verbs were replaced by synonyms (e.g. "The spy ripped up the documents"). The task was to identify the originally read sentences. Performance on the test indicated that the exact wordings of the sentences were present in memory. Originally read sentences were recognized with an average confidence rating of 2.14 (the maximum possible rating was 3), while synonymous wordings of the sentences were recognized with an average confidence rating of only .92. If the sentences had been stored in canonical meaning representations in memory, originally read and synonymous sentences should have been recognized with equal confidence. More confident recognition of originally read sentences suggests that the surface features of sentences, such as particular wordings, were retained as functional componenets of their memory representations.

In the second experiment, people again studied thirty-two sentences for meaning. Later, they were given a verification test. The test included two kinds of "true" sentences: the originally studied sentences and those same sentences with the verbs replaced by synonyms. It also included an equal number of "false" sentences, constructed by replacing the verbs in studied sentences with unrelated verbs. The task was to respond "TRUE" to the true sentences and "FALSE" to the false sentences. Performance indicated that the exact wordings of the sentences were present in memory. Reaction time to verify the originally studied sentences was 1.14 sec, while reaction time to verify the

synonymous sentences was 1.41 sec. If the sentences had been stored in canonical meaning representations in memory, originally read and synonymous sentences should have been verified equally fast. Faster verification of the originally studied sentences suggests that the words in the sentences were retained as functional components of their memory representations.

In the third experiment, people studied twenty-four sentences for meaning. Eight of the sentences were unrelated to other studied sentences. Eight sentences shared a common verb with other sentences. Eight of the sentences contained verbs whose synonyms appeared in other sentences. Later, a verification task was given, in which people had to respond "TRUE" to the studied sentences and "FALSE" to an equal number of sentences whose verbs had been replaced by unrelated verbs. There was no difference between reaction times to verify sentences in the unrelated and synonymous verbs conditions (.994 sec versus 1.006 sec). Reaction time to verify sentences in the common verbs condition, however, was significantly longer (1.092 sec). These results indicate that interference in memory for related sentences occurs when sentences share explicit surface information, not when they share only semantic information. It is generally assumed (cf. [3, 41]) that interference in sentence memory arises from shared representational structures. When sentences share a common word, interference presumably derives from a shared representation of that word. If words and sentences have canonical meaning representations in memory, synonymous

words in learned sentences should also produce shared representational structures and equivalent interference effects. The absence of interference effects among sentences sharing synonymous words suggests that distinctive surface information, rather than equivalent semantic information, is the basis for sentence representation in memory.

The results reviewed above rule out any typical theory of canonical meaning representation. Instead, they favor a more literal representational scheme, such as an integrated semantic-propositional network. Of course, the results do not rule out the coexistence of canonical meaning representations and corresponding representations of surface information. However, such a combined theory would be a less parsimonious description of human pattern matching, and such double representations would be considerably less efficient than either one alone for artificial knowledge systems.

Literal representations are, in some ways, more efficient than canonical representations. Because canonical representations minimize the number of representational component types, they maximize the number of component tokens and relations among tokens necessary to represent a given proposition. Thus, by employing a more informative code, literal representations require less storage than canonical meaning representations. Pattern matching should also be more efficient with literal representations than with canonical

representations because they are less complex and, therefore, require fewer basic matching operations. Interference in pattern matching should be minimized with literal representations because, by increasing the number of distinctive symbols used to construct meaning representations, they minimize the number of partial matches for a given target pattern. Finally, literal representations support inferencing mechanisms that are as efficient as those based on canonical representations. For example, it is possible to express general inferences as productions. For the case described above, such a rule might be [x is the agent of verb w and w is an element of {"run," "walk," "go"...}] => (there exist x, y, z): x went from y to z. Thus, in addition to being a better description of human pattern processing, literal meaning representations may, in fact provide a more efficient data base for artificial knowledge systems.

Prototypes as Representations of Classes of Stored Patterns

Many distinctive patterns represent variations on prototypic patterns. For example, there are numerous discriminable instances of the pattern "tree" or the pattern "dinner at a restaurant." Many designers of artificial knowledge systems (c.f. [8]) propose that pattern classes should be represented by single prototypes (frames, schemas, or scripts), rather than by representations of all individual instances of the pattern. The prototype emphasizes

the similarities and obscures the peculiarities among instances. Prototypes apparently provide efficiencies in storage, matching, and inferencing. Because the information in a class of patterns can be concentrated in a single prototype, the amount of storage required is minimized. Similarly, the set of patterns that must be searched to match a class member is reduced from the set of all previously encountered class members to the single abstracted prototype. In addition, pattern-directed inferencing is facilitated, because the system can use the "typical" information present in the prototype to fill in any gaps in the information present in a target pattern.

Several kinds of empirical support have been adduced in support of the assumption that people represent classes of patterns as prototypes. For example, several studies have shown that when people experience many variations on prototypic patterns, their subsequent recognition and classification of both familiar and unfamiliar instances of the patterns are strongly influenced by the degree of similarity between the test pattern and the prototype [17, 31, 32]. The more similar an instance is to the prototype, the better it is recognized and classified. The prototype is recognized and classified best of all, even if it has never been experienced previously. Other studies [29, 40] have shown that peoples' memory for texts that they have read can be predicted on the basis of prototypic text structures; higher-order propositions in the prototypic

structure are remembered better than lower-order propositions.

Despite the consistency of these psychological results with prototype theory, a recent study [20] has shown that people do not represent and process all classes of patterns in terms or their relationships to prototypes. In this study, people studied descriptions of individuals (name, age, education, marital status, hobby) who belonged to one of three classes. Each individual description was a variation on a prototypic description for the appropriate class. The goal for people in the experiment was to learn to classify individuals correctly. After studying 132 individual descriptions, they made recognition judgments (have I ever seen this particular description before?) and classification judgments (in what class does this description belong?) about familiar and unfamiliar descriptions.

The prototype theory assumes that during study, people abstracted the most typical features of the individual descriptions in each class and stored them as prototypic class members. Subsequent recognition and classification judgments would be based presumably on the similarity of a test description to each of the stored prototypes. An alternative theory (property-set theory) assumes that during study, people stored and strengthened all features and combinations of features that they encountered in individual descriptions, along with appropriate class designations.

Subsequent recognition and classification judgments would be based presumably on the stored strengths of the features and combinations of features in the test description. The two theories make a variety of differential predictions regarding recognition and classification judgments. For example, the prototype theory predicts that descriptions of prototypes should be recognized with greater confidence and classified with greater accuracy and confidence than any other descriptions. The property-set theory, on the other hand predicts that descriptions should be recognized better and classified better and with higher confidence than the prototype whenever their constituent features and combinations of features occur more often than those in the prototype.

The prototype and property-set theories were comparatively evaluated in terms of their ability to predict relative recognition judgments for twenty-four critical pairs of descriptions (i.e. pairs for which the two theories made different predictions). These included six comparisons of prototypes to non-prototypic descriptions. For these, the prototype theory predicted more confident recognition for the prototype, while the property-set theory predicted better recognition confidence for the non-prototypic description. They also included twelve comparisons of non-prototypic descriptions that were equidistant from the prototype. For these, the prototype theory predicted equally confident recognition for paired descriptions, while the property-set theory predicted that

one of the two should be recognized more confidently. The remaining six comparisons were between pairs of descriptions that were different distances from the prototype. For these, the prototype theory predicted better recognition of the descriptions that were nearer the prototype, while the property-set theory made the opposite prediction. The property-set theory fit the data considerably better than the prototype theory (twenty-three of the twenty-four critical predictions confirmed). The two theories were also comparatively evaluated in terms of their ability to predict relative recognition and classification judgments for all pair-wise combinations of twenty-eight test descriptions. These included pairs of items varying in both distance from prototypes and frequencies of constituent feature combinations. The property-set theory provided a significantly better fit to both the recognition and classification data than the prototype theory.

These results indicate that people do not restrict the information stored in memory to those features that occur most often among a class of patterns, as assumed by the prototype theory. Rather, people appear to store all of the information present in the patterns they experience and they use all of it during pattern matching tasks such as recognition and classification.

In addition to modeling human pattern processing more accurately, representations that preserve all of the information

in a class of patterns permit certain intelligent behavior that would not be possible if the data base were constrained to preserve only prototypic information. For example, there may be information among the instances that is not frequent enough to appear in a prototype (especially information that represents an alternative to information present in the prototype), but which is, nonetheless, diagnostic. This information will be available for appropriate use if all information in the class of patterns is stored, but not if only prototypes are stored. Also, it may not be possible to anticipate the importance of alternative values on prototypic dimensions in dealing with future situations. If only prototypes are stored, the system must make do with arbitrarily selected values. If all the information in the class of patterns is stored, the system has complete information available to be used as needed.

Processing Partially Matched Patterns

In any reasonably sophisticated knowledge system, many stored patterns share common subpatterns. For example, a system that knows about U.S. history is likely to have many stored propositions about World War II. Each of these propositional representations will include a subrepresentation of the concept "World War II." When the system attempts to match a particular target proposition to its memory representation, the presence of other propositional representations with common

subrepresentations will produce many partial matches. The pattern matching process will be slowed by evaluation of all representations in memory that partially match the target proposition. This kind of interference occurs in several artificial pattern matching systems [2, 8, 23].

Interference is also one of the best known phenomena in human learning and memory. For example, several studies [3, 19, 21, 28, 41] have shown that reaction time to verify a stored proposition is an increasing function of the number of learned propositions that share concepts with the test proposition. Presumably, all learned propositions that include concepts present in the test proposition produce partial matches and, therefore, must be evaluated during the verification process.

While interference from partial matches is common, it is not inevitable. In a recent study [18], it has been shown that interference occurs only when the target pattern is not well known. This seems intuitively reasonable when one considers, for example, that although most people know many more propositions about themselves than they do about anyone else, it is precisely those propositions that are the most accessible. In this study, people learned propositions that shared common concepts with zero, one, two, or five other learned propositions. They then had to verify (classify as true or false) the propositions they had studied, along with some false propositions. When one of the propositions

was presented for verification, any other learned propositions with common concepts presumably produced partial matches to the target. These partial matches should have interfered with (slowed) the verification process. In fact, when people had learned the propositions well enough to verify them accurately (<5% errors), reaction time was an increasing function of the number of partially matched propositions in memory. However, as the amount of experience people had at verifying the propositions increased, interference decreased, ultimately to zero. That is, after a reasonable amount of practice (about one hundred trials), reaction time was constant, regardless of the number of partially matching patterns in memory. Figure 1 illustrates these results. These data suggest that while it may be necessary to search all partial matches for complete desired matches to some learned patterns, it is not necessary for well learned patterns.

The ability to bypass partial matches in the search for a desired complete match should make pattern matching faster and more accurate. This ability is essential in a system that must acquire more than a few propositions regarding any one topic. It is generally true that the more important a certain topic is to a person or machine, the more propositions regarding that topic must be stored in memory. It is undesirable to have a knowledge system in which pattern matching gets slower as the number of related propositions increases. This would be tantamount to a system in which matching gets slower as the

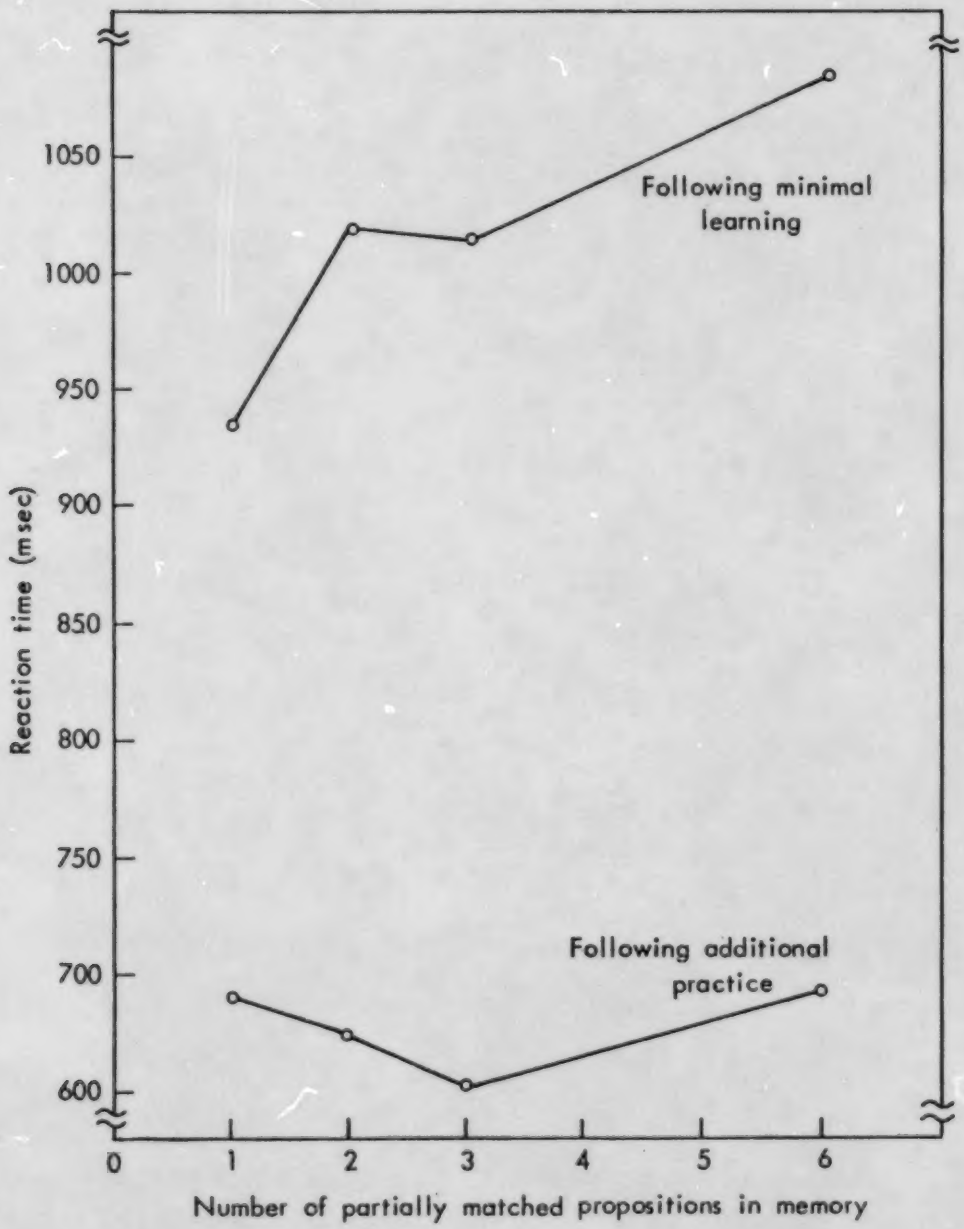

Figure 1: Effect of partially matching propositions in memory on verification of target propositions.

importance of the knowledge increases. People apparently respond to such conditions by making qualitative changes in the pattern matching process, such that the process is no longer slowed down by the existence of partially matching patterns in memory. A comparable capacity in machine systems is an obvious goal of future research.

Effect of Pattern Complexity and Capacity Limitations

All artificial pattern matching systems operate under capacity limitations. A system has only finite resources. If pattern matching operations are performed serially, the time required to complete a match must be directly related to the number of operations performed. The same must also be true of parallel matching processes once the number of requested simultaneous operations performed exceeds the system capacity. The number of operations performed is, in turn, determined by the complexity of the pattern to be matched, i.e. the number of elements and relations necessary to specify the pattern. For example, a human face is a more complex pattern than an equilateral triangle and, therefore, should take longer to match. In fact, it is widely believed [1, 22, 26] that most interesting problems involving recognition, matching, and searching of general graph structures are intractable; the amount of computing needed for solution is an exponential function of the number of elements in the graph structures compared.

Similarly, it is a common assumption that because people

have limited processing capacity, increasing the complexity of a pattern must decrease peoples' facility at processing it. This relationship is illustrated, for example, in the finding that the more elements in a proposition to be verified, the longer the verification time [14]. Other studies [11, 19] have shown that the time required to verify a semantic or relational proposition is an increasing function of the number of elements of the memory structure that must be activated (see also [12, 36]).

On the other hand, some researchers have shown that complex patterns are sometimes processed as well as or better than simpler patterns. For example, several studies [10, 13, 33] have shown that chess and Go masters operate on complex board configurations at least as well as non-masters operate on simple configurations. These results suggest that, with experience, chess and Go players acquire the ability to process increasingly complex patterns without any decrement in performance. Other studies have shown that people can process some words and phrases as well as or better than they can process their constituent letters or words [7, 24, 27].

A recent paper [18] has argued that increasing experience can enable people to process complex patterns more easily than they process simpler patterns. In an earlier study [14], people studied subject-verb-object-location-time propositions. After they had learned the sentences minimally well (after seven

trials), they verified two-part or three-part subsets of the propositions (e.g. subject-verb, subject-verb-object) along with the complete propositions. It was found that people verified two part subsets five msec faster than three-part subsets and verified three-part subsets twenty-six msec faster than complete (five-part) propositions. The later study [18] involved a similar procedure using subject-verb-object propositions. Here, when people had learned the complete propositions moderately well (after eighteen trials), they verified one-part and two-part subsets equally fast. However, after they had gained considerably more experience with the complete propositions (after one hundred trials), they verified two-part subsets 143 msec faster than one-part subsets. Thus, with minimal learning, simple patterns were matched faster than complex patterns, while with considerably better learning, complex patterns were matched faster than simpler patterns.

These results suggest that when propositions are moderately well learned, they are represented and processed as patterns of elements and relations. Matching a sub-proposition requires matching each of its substituent elements and relations. Because matching time is determined by the number of elements and relations that must be matched, simple propositions are matched faster than complex propositions. When propositions are very well learned, however, they are represented and processed as "unitary" events. Matching a sub-proposition requires matching and then "decomposing" the corresponding complete proposition. In

addition, matching time for the complete proposition is assumed to be determined by the completeness of the target proposition. Thus, complex (more complete) patterns are matched faster than simpler patterns.

The acquired ability to match complex patterns as quickly and as accurately as simple patterns is a desirable goal for artificial knowledge systems too. A truly adaptive knowledge system should be sensitive to increases in the importance of certain stored patterns, as evidenced, for example, by frequency of access. People apparently respond to such conditions with qualitative changes in pattern representation and processing that compensate for apparent capacity limitations. Similarly, if artificial knowledge systems are to become as effective in their task domains as people are in theirs, effective means must be developed to increase their capacity to process important, but complex, structures.

Conclusions

Human cognition can provide an inspiring model for the design of artificial knowledge systems. This paper has presented evidence that four design principles commonly embodied in artificial knowledge systems are inconsistent with human cognitive capabilities. To summarize: Human memory representations appear to be knowledge extensive, preserving the surface features of each stored pattern, including those

that are members of specifiable classes of patterns. Human pattern matching processes appear to be both powerful and flexible, overcoming apparent capacity limitations and bypassing the evaluation of partial matches in order to expedite the processing of important patterns. While such capabilities would be advantageous in intelligent systems, they are unfortunately beyond the state of the art in artificial intelligence. Memory and capacity limitations appear inevitable, and no search algorithm is yet known that can avoid extensive evaluation of all partial matches. While it is always possible to design functionally equivalent systems using alternative mechanisms, it has proven extremely difficult to discover any mechanism capable of producing anything approaching the intelligent and adaptive behavior typical of an average human being. Although several successful special-purpose artificial knowledge systems have been developed, the limitations are as salient as their capabilities. In the long run, the most promising avenue of approach toward the development of artificial knowledge systems may be to investigate and model human cognition as closely as possible. The four properties of human pattern processing described in this paper delimit a great gap between the technologies of natural and artificial intelligence. They illustrate both the difficulty and the potential rewards of modeling artificial intelligence systems after the human knowledge system.

References

1. Aho, A.V., Hopcroft, J.E., and Ullman, J.D. The design and analysis of computer algorithms. Addison-Wesley, Reading, Mass., 1974.

2. Anderson, J.R. Language, memory, and thought. Lawrence Erlbaum Associates, Murray Hill, N.J., 1976.

3. Anderson, J.R. Retrieval of propositional information from long-term memory. Cognitive Psychology, 6, 1974, 451-474.

4. Anderson, R.C. Encoding processes in the storage and retrieval of sentences, Journal of Experimental Psychology, 2, 1971, 338-340.

5. Anisfeld, M. False recognition produced by semantic and phonetic relations under two presentation rates. Psychonomic Science, 17, 1969, 366-367.

6. Anisfeld, M., and Knapp, M. Association, synonymity, and directionality in false recognition. Journal of Experimental Psychology, 77, 1968, 171-179.

7. Bower, G.H. Chunks as interference units in free recall. Journal of Verbal Learning and Verbal Behavior, 8, 1969, 610-613.

8. Bobrow, H.G., and Winograd, T. A knowledge representation language. Cognitive Science, in press.

9. Brewer, W.F. Memory for ideas: Synonym substitution. Memory and Cognition, 3, 1975, 458-464.

10. Chase, W.G., and Simon, H.A. Perception in chess. Cognitive Psychology, 4, 1973, 55-81.

11. Collins, A.M., and Quillian, M.R. Retrieval time from semantic memory. Journal of Verbal Learning and Verbal Behavior, 8, 1969, 240-247.

12. Crowder, R.G. Behavioral strategies in immediate memory. Journal of Verbal Learning and Verbal Behavior, 8, 1969, 524-528.

13. DeGroot, A.D. Perception and memory versus thought: Some ideas and recent findings. In Kleinmuntz, B. (Ed.), Problem Solving. Wiley, New York, 1966.

14. Dosher, B.A. The retrieval of sentences from memory: A

speed-accuracy study. Cognitive Psychology, 8, 1976, 291-311.

15. Fillenbaum, S. Memory for gist: Some relevant variables. Language and Speech, 9, 1966, 217-227.

16. Fillenbaum, S. Words and feature complexes: False recognition of antonyms and synonyms. Journal of Experimental Psychology, 22, 1969, 400-402.

17. Franks, J.J., and Bransford, J.D. Abstraction of visual patterns. Journal of Experimental Psychology, 90, 1971, 65-74.

18. Hayes-Roth, B. The evolution of cognitive structures and processes. Psychological Review, 1977, in press.

19. Hayes-Roth, B., and Hayes-Roth, F. Plasticity in memorial networks. Journal of Verbal Learning and Verbal Behavior, 14, 1975, 506-522.

20. Hayes-Roth, B., and Hayes-Roth, F. Concept learning and the recognition and classification of exemplars. Journal of Verbal Learning and Verbal Behavior, in press.

21. Hayes-Roth, B., and Hayes-Roth, F. The prominence of lexical information in memory representations of meaning. Journal of Verbal Learning and Verbal Behavior, 1977, in press.

22. Hayes-Roth, F. Representation of structured patterns and efficient procedures for their recognition. Pattern Recognition, 1976.

23. Hayes-Roth, F., Mostow, D.J., and Fox, M.S. Understanding speech with the Hearsay-II system. In Bolc, L. (Ed.) Natural language communication with computers. Springer-Verlag, Berlin, 1977.

24. Horowitz, L.M., Day, R.S., Light, L.L., and White, M.A. Availability growth and latent verbal learning. Journal of General Psychology, 78, 1968, 65-83.

25. Jackendoff, R. A system of semantic primitives. In Theoretical issues in natural language processing, Proceedings of an Interdisciplinary Workshop in Computational Linguistics and Artificial Intelligence, Cambridge, Massachusetts, 1975, pp. 24-29.

26. Karp, R.M. Reducibility among combinatorial problems. In Miller, A.E. and Thatcher, J.W. (Eds.), Complexity of computer computations, Plenum Press, New York, 1972.

27. LaBerge, D., and Samuels, S.J. Toward a theory of automatic information processing in reading. Cognitive Psychology, 6, 1974, 293-323.

28. Lewis, C.H., and Anderson, J.R. Interference with real world knowledge. Cognitive Psychology, 8, 1976, 311-335.

29. Mandler, J.M., and Johnson, N.S. Remembrance of things parsed: Story structure and recall. Cognitive Psychology, 9, 1977, 111-151.

30. Miller, G.A. English verbs of motion: A case study in semantics and lexical memory. In Melton, A.W. and Martin, E. (Eds.), Coding processes in human memory, John Wiley and Sons, New York, 1972.

31. Oldfield, R.C. Memory mechanisms and the theory of schemata. British Journal of Psychology, 56, 1965, 349-358.

32. Posner, M.I. Abstraction and the process of recognition. In Bower, G.H. and Spence, J.T. (Eds.), Psychology of learning and motivation, Vol. 3. Academic Press, New York, 1969.

33. Reitman, J.S. Skilled perception in Go: Deducing memory structures from inter-response times. Cognitive Psychology, 8, 1976, 336-356.

34. Rumelhart, D.E. Notes on a schema for stories. In D.G. Bobrow and A. Collins (Eds.), Representation and understanding: Studies in cognitive science. Academic Press, New York, 1975.

35. Sachs, J.S. Memory in reading and listening to discourse. Memory and cognition, 2, 1974, 95-100.

36. Savin, H.B., and Perchonock, E. Grammatical structure and the immediate recall of English sentences. Journal of Verbal Learning and Verbal Behavior, 4, 1965, 348-353.

37. Schank, R.C. Conceptual dependency: A theory of natural language understanding. Cognitive Psychology, 3, 1972, 552-631.

38. Schank, R.C., Goldman, N., Rieger III, C.J., and Riesbeck, C. Margie: Memory, analysis, response generation, and inference on English. Proceedings of the Third International Joint Conference on Artificial Intelligence, 1973, 255-261.

39. Schank, R.C., and Rieger III, C.J. Inference and the computer understanding of natural language. Artificial Intelligence, 5, 1974, 373-412.

40. Thorndyke, P.W. Cognitive structures in comprehension and memory of narrative discourse. Cognitive Psychology, 9, 1977, 77-110.

41. Thorndyke, P.W., and Bower, G.H. Storage and retrieval processes in sentence memory. Cognitive Psychology, 6, 1974, 515-543.

42. Wilks, Y. Grammar, meaning, and the machine analysis of language. Routledge and Kegan Paul, Ltd., London, 1972.